Digital Bas-Relief from 3D Scenes

Tim Weyrich*   Jia Deng*   Connelly Barnes*   Szymon Rusinkiewicz*   Adam Finkelstein*

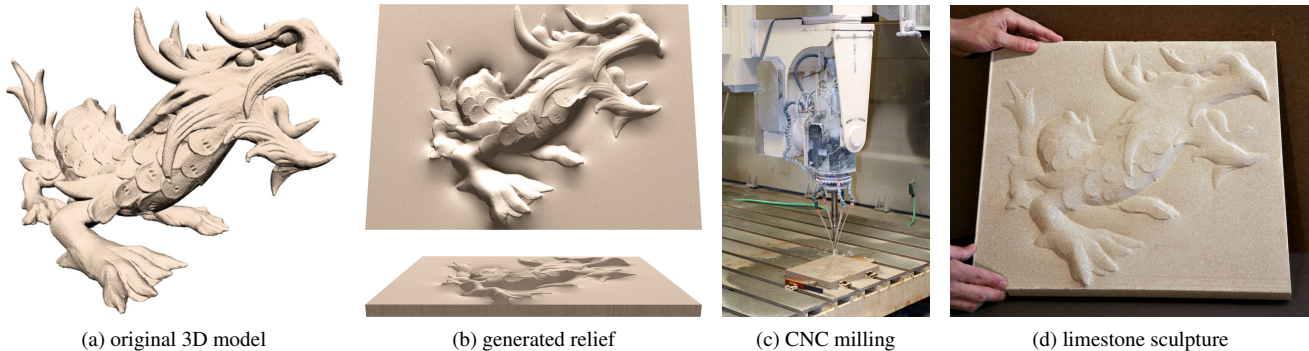Princeton University

| (a) original 3D model | (b) generated relief | (c) CNC milling | (d) limestone sculpture |

Figure 1: *Creating a bas relief sculpture. (a) Starting with a 3D model as well as a few parameters that adjust aesthetic qualities, (b) our system creates a height field visually matching the input from a particular vantage point – shown above – but that incorporates a dramatically reduced range of heights – below. (c) The height-field is cut on a computer-controlled milling machine to produce (d) a physical bas relief sculpture carved in limestone, roughly 15" tall.*

Abstract

We present a system for semi-automatic creation of bas-relief sculpture. As an artistic medium, relief spans the continuum between 2D drawing or painting and full 3D sculpture. Bas-relief (or low relief) presents the unique challenge of squeezing shapes into a nearly-flat surface while maintaining as much as possible the perception of the full 3D scene. Our solution to this problem adapts methods from the tone-mapping literature, which addresses the similar problem of squeezing a high dynamic range image into the (low) dynamic range available on typical display devices. However, the bas-relief medium imposes its own unique set of requirements, such as maintaining small, fixed-size depth discontinuities. Given a 3D model, camera, and a few parameters describing the relative attenuation of different frequencies in the shape, our system creates a relief that gives the illusion of the 3D shape from a given vantage point while conforming to a greatly compressed height.

Keywords: sculpture, non-photorealistic rendering, tone mapping, geometry generation

1 Introduction

Relief is a method of sculpture in which forms are carved into a relatively flat surface, and in essence creates a bridge between full 3D sculpture and 2D media such as painting. On this spectrum *alto-relievo* (high relief) is closer to full 3D, whereas flatter works are described as *basso-relievo* (low relief, and more conventionally called *bas-relief*). Examples of bas-relief may be found throughout history (Figure 2) and excellent examples survive from many ancient cultures including Greek, Persian, Egyptian, Mayan, and

---

*Email: [tweyrich,jiadeng,csbarnes,smr,af]@cs.princeton.edu

Indian art and architecture [Flaxman 1829; Hoffman 1939; Read 1961]. Today bas-reliefs are commonly found in a variety of media in architecture, industrial design and coins. However, even with the advent of 3D modeling, computer-driven milling equipment, and 3D printing technology, the *design* of bas-relief sculptures remains largely in the hands of artists.

This paper presents a tool to assist artists in designing bas-reliefs based on 3D models. The designer starts with (or constructs) a 3D model, together with a camera view. Our system creates a bas-relief based on the model, camera, and selection of a few intuitive artistic controls such as the final output height and parameters impacting relative attenuation of large- or fine-scale detail. The generated bas-relief respects the constraint on final height while attempting to simulate for orthogonal views the effect of the original model as seen from the chosen camera. Nevertheless, as the viewer moves to an off-orthogonal vantage point, it becomes apparent that the model

Figure 2: Top: *Ancient Greek, Assyrian relief.* Below: *Modern examples.*

is distorted and generally flat (Figure 1(b), bottom). This effect relies on a feature of the human perception known as the "bas-relief ambiguity" which has been recently studied in the computer vision community [Belhumeur et al. 1999] but known to artists for centuries.

The technique described herein relies on several observations. From the point of view of the camera, the model may be described as a height field, with two implications. First, rather than considering all possible model deformations in 3D, we can limit ourselves to the space of transformations on a 2D height field. Second, by sampling the heights on a 2D grid, and recognizing that our goal is generally to reduce the dynamic range of these heights, we can tap into the broad literature for "tone mapping", whose goal is to reduce the dynamic range in high dynamic range (HDR) images. Finally, unlike HDR imagery, the resulting relief must be continuous (in order not to present any gaps or vertical walls in the height field, apparent from slightly off orthogonal views). Therefore, our method must collapse empty space at object silhouettes, essentially "squashing" foreground objects down against the background objects that they occlude. Fortunately, this requirement supports the goal of reducing dynamic range in the resulting relief.

For effective bas-relief there is a strong tension between the goal of preserving the appearance of the shape – largely driven by perceived gradients and curvatures – and the requirements that the shape be continuous and greatly flattened. Consider the inset figure wherein two crossed tubes A and B lie over flat background C. A proper bas-relief of this scene requires that the transitions labeled AB, AC and BC all have locally similar properties, conflicting with the goal of globally preserving the shapes and depth ordering. This paper focuses on bas-relief, as opposed to medium or high relief, because of the peculiar challenges inherent in this form. Nevertheless, the solutions we present may also assist in the creation of high relief where these tensions are not as strong. Hereafter in this paper we will generally use "relief" to mean "bas-relief."
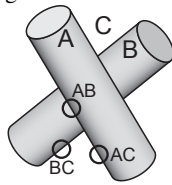
Note that the goal of this work is not to replace the artist, but rather to provide high level tools for creating bas-relief. The role of the artist remains, as always, one of composition and the tradeoff between abstraction and representation. Artists have for centuries used mechanical devices to assist in the design of sculpture (e.g. the pantograph [Hoffman 1939]) and paintings (e.g. the camera lucida and camera obscura [Hockney 2001]). Some modern sculptors use 3D modeling software, and some use automatic milling devices [Späth 2006], and the tools we describe would be advantageous in such settings. The applications for this work apply in any setting where shapes are designed to be carved into wood or stone, fabricated on a 3D printer, injection molded, embossed into a product package, or stamped from a die. The latter category is broad, including objects ranging from coins to greeting and business cards which often have low relief stamped from a press.

1.1 Bas-Relief

Here we briefly review the principles that artists have developed and codified over many centuries for creation of effective bas-relief [Flaxman 1829; Hoffman 1939; Read 1961]. These principles motivate the algorithmic decisions described hereafter.

Depth. The illusion of depth is mainly achieved by perspective foreshortening in the image plane.

Object order. The relative depth order of overlapping objects must be preserved. However, non-interacting objects may be moved to different heights, relying on foreshortening for depth cues.

Compression. In many cases, objects in the the background are more compressed (flatter) than those in the foreground.

Discontinuities. Depth discontinuities in the scene *must not* give rise to depth discontinuities in the relief. An extension from the foreground object to the background object would look wrong when viewed even slightly from the side.[1]

Steps. Depth discontinuities in the scene *may* give rise to a small step that traces along the object outline. The step height should be constant across the relief.[2]

Undercuts. In some materials, the artist may add a small "undercut" to these steps in order to accentuate silhouettes.

Materials. Reliefs have been carved into a vast array of materials including wood, clay, stucco, metal, stone, ivory, bone, ice, and now industrial product design, e.g., by injection molding. Each material exerts its own constraints on the problem. For example, it is difficult to carve very fine detail into marble, and it would be difficult to see due to subsurface scattering. Another example is that of coins, which are stamped by a die and hence cannot have undercuts; furthermore, modern coins only have shallow relief because they have to be stackable.

2   Related Work

As mentioned in Section 1, by treating the 3D model as a height field from the point of view of the camera, we are able to borrow some methods developed for tone mapping of HDR images. Our goal is to "squeeze" these heights into the dynamic range available in our relatively flat medium, while preserving the original look. Of course, the look of a shape in 3D does not follow the same principles of perception as that of a grid of pixel intensities, so some of the tone mapping literature as pertains to perceptual or photographic processes (e.g., [Larson et al. 1997; Reinhard et al. 2002; Tumblin and Turk 1999]) would be difficult to apply. However, general methods for dynamic range compression may be adapted. DiCarlo and Wandell [2000] offer a review of the basic methods for tone mapping, dividing the literature into two general approaches that employ either a *tone reproduction curve* (TRC) or a *tone reproduction operator* (TRO).

TRC-based methods provide for the entire image a single monotonic curve that globally maps input intensities to output intensities. An example is the histogram adjustment approach of Larson et al. [1997]. TRC-based methods will not make good use of the dynamic range available for the bas-relief when there are objects distributed throughout the range of heights in the scene (even if all the high values are on the left and the low values are on the right). Moreover, there is no way to use a monotonic remapping of the heights to eliminate height discontinuities.

Therefore, we turn our attention to TRO-based schemes, whose general framework is to use varying tone mapping throughout the image so that locally the remapping is like a TRC-method (but globally the mapping varies). Several TRO schemes use some form of multi-scale transformation on the input, separating high-frequency detail from the base shape, attenuating the base shape, and then adding back in the detail [Tumblin and Turk 1999; Ashikhmin 2002; Durand and Dorsey 2002]. Our system implements similar controls, as they allow the artist to specify the relative attenuation of fine detail versus coarser shapes. The challenge in these methods is to do the decomposition in such a way as to avoid ringing artifacts at sharp features due to high-frequency information being encoded in the (attenuated) low-frequency shapes. Durand and Dorsey [2002] use bilateral filters to ameliorate this problem, whereas Tumblin and Turk [1999] divide the scene into smooth regions bounded by discontinuities. In our problem, we start with the advantage of know-

---

[1]This principle does need not apply in high relief, where a foreground object may be fully formed in 3D and stand out from the background.

[2]An exception is where the relief climbs the façade of a large building and larger steps may be used higher up.
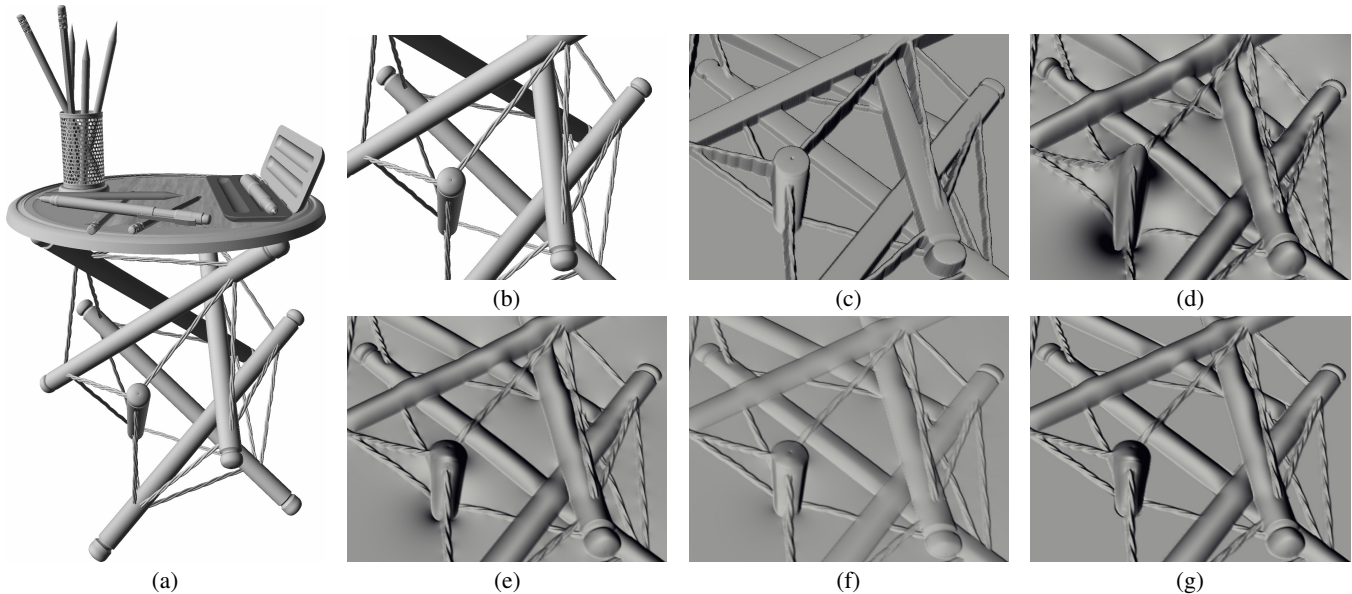
Figure 3: *Different depth compressions applied to the same input scene.* (a) *Input scene.* (b) *Close up of this scene.* (c) *Linear scaling of perspective depth values after Cignoni et al. [1997].* (d) *Relief generated from our technique, using no depth-compression.* (e) *Depth compression for* $\alpha = 1$. (f) *Compression for* $\alpha = 10$. (g) *Same as (e), with boundary conditions at silhouettes.*

ing from the 3D scene exactly where the sharp discontinuities lie (silhouettes), so we implement a diffusion-based scheme respecting the silhouettes in order to extract low-frequency shapes. Finally, a recent TRO-based scheme devised by Fattal et al. [2002] works in the gradient domain to attenuate high-gradient areas more than low-gradient areas. Briefly, this scheme converts the intensities of an HDR image into the gradient domain, and attenuates the high magnitude gradients more than others. The resulting field is non-integrable, but a least squares best match is used to integrate back, producing an image with reduced dynamic range. A multiresolution scheme is used to solve the least-squares system efficiently. We have adapted this method for bas-relief, allowing us to explicitly bound the gradient everywhere – crucial at silhouette boundaries where otherwise we might leave a discontinuity.

This paper may generally be thought of as falling under the umbrella of non-photorealistic rendering (NPR). The system presented here simulates the effect of an artistic medium, like much of the work in NPR which has, starting from a 3D model, simulated renderings in various traditional media such as oil painting [Meier 1996], pen and ink [Winkenbach and Salesin 1996], and technical illustration [Gooch et al. 1998]. However, the majority of this work has focused on creating 2D imagery, and it is unusual to find other examples of 3D models transformed into "renderings" that are meant to be viewed in 3D. One exception is the system of Raskar et al. [2002], in which dioramas depict motion through the use of projected light. Also, Sourin [2001] developed a system to allow artists to perform virtual embossing and woodcuts, producing low reliefs in 3D using virtual analogs of traditional sculpting tools entirely "by hand." Finally, the branch of NPR that focuses on making line drawings from 3D models (e.g., [Markosian et al. 1997; Hertzmann and Zorin 2000; DeCarlo et al. 2003]) bears relation in that we extract linear features (silhouettes) and "draw" them as steps or undercuts.

One paper to date has directly addressed the problem of semiautomatic creation of relief sculptures [Cignoni et al. 1997]. This work makes several observations on which our technique relies. First, it notes that the problem can be solved over a height field from the point of view of the camera; and second, that by capturing those heights as *z*-buffer values after camera transformation, a nonlinear scaling is given to the heights that gives more dynamic range to the objects close to the camera than those that are far, while preserving planarity in the result. The final step in their algorithm is

simply to linearly scale these heights to fit within the available dynamic range. This method works well for high relief and for scenes of limited depth complexity. However, in the sense that it falls under the umbrella of TRC-based schemes, it will not work well for bas-relief in general, considering the aforementioned requirements for dynamic range compression and continuity.

3 Bas-Relief from 3D Models

Our method begins with a general polygonal 3D scene as input and transforms it into a regularly sampled height field (Section 3.1). In a depth compression stage we adapt the height field to meet the requirements of bas-relief (Section 3.2). This compression takes place in the gradient domain and is followed by an integration step that recovers a height field of the resulting relief from the modified gradients (Section 3.3).

3.1 Scene Input

Cignoni et al. [1997] observe that retrieving the depth values of a scene from a perspectively mapped depth buffer results in desirable properties, such as depth foreshortening proportional to the perspective foreshortening in the image plane, and the preservation of planarity of surfaces, which is desirable, e.g., for architectural models. Following this observation, we start from a perspective depth buffer of an arbitrary 3D scene. Such a depth buffer mapping is, for example, provided by any OpenGL viewer by default. OpenGL's depth buffer convention perspectively maps z values between the near and far plane to depth values within $[0, 1]$. Choosing the near and the far plane to tightly enclose the scene accordingly delivers a plane-preserving height field of the input scene with a natural foreshortening of distant features, see Figure 3(a–c). While this kind of foreshortening is one of the central challenges for a sculptor, homogeneous geometry provides it for free.

3.2 Depth Compression

While Cignoni et al. [1997] consider plane preservation the topmost goal, we find that this constraint unduly hinders an effective, shape-preserving depth compression: in order to meet additional constraints such as the continuity at silhouettes, the height field would have to be scaled down to a degree that flattens out important features of the scene (see Figure 3(c)). Instead, we loosen the planarity constraint in favor of a more perceptually-based compres-

sion scheme. The most relevant perceptual features that have to be preserved during the transition from a scene to its bas-relief are silhouettes at depth-discontinuities and shading under incident illumination. Both aspects are central visual cues that convey shape.

Silhouettes are mapped to minute steps in the height field that accentuate transitions between neighboring surface patches in the continuous relief. In order to preserve shading cues, we try to preserve surfaces gradients of the input height field $h(x,y)$ during the shape transformation.

Here we provide a brief overview of our approach to achieving these goals; it is discussed in greater detail through the remainder of the section. We initially differentiate the height field $h(x,y)$, performing the depth compression in the gradient domain. In particular, we fix the gradient directions $\mathbf{v}(x,y) = \nabla h(x,y)/\|\nabla h(x,y)\|$, confining the subsequent operations to the gradient magnitudes. Our depth compression scheme attenuates large gradients that correspond to large slopes in the height field, effectively attenuating steep scene elements. After gradient manipulation, the heights $h'(x,y)$ of the bas-relief are recovered by integrating over the modified gradient $g'(x,y)$. As an optional step, we propose a simple frequency-based approach to control feature scale in the final relief.

The rationale behind this general approach is that it attempts to preserve gradient directions while inducing monotonic and smoothly-varying slopes in the gradient magnitude. Because the gradient is orthogonal to *isophotes* (lines of equal intensity) under headlight illumination, this approach tends to preserve shapes visible in the lit image.

Silhouettes. Internal and exterior silhouettes require special treatment in the gradient field, but allow for drastic depth compression. In our pipeline we use forward differences to compute gradients. Hence, depth discontinuities at silhouettes within the scene contribute as overly large gradients that would give rise to depth discontinuities when integrating back into the height domain. In order to cause the final integration step to map surface patches next to a silhouette to the same height, we set gradient values at silhouettes to zero. Silhouettes are detected by thresholding the gradient magnitudes: a height field pixel is treated as part of a silhouette if $\|\nabla h(x,y)\| > \vartheta_{\text{sil}}$. On coarsely sampled input height fields, this can lead to an erroneous classification of high slopes as silhouettes; however, for higher sampling resolution, silhouette gradients exceed other scene gradients by an order of magnitude.

Gradient Compression. Our depth compression applies a nonlinear compression function C to the gradient magnitude, largely preserving small gradients while attenuating large slopes. We found that

$$C(x) = \frac{1}{\alpha}\log(1+\alpha x), \quad \alpha > 0 \, , \tag{1}$$

provides reasonable compression rates without overly distorting shapes. Figure 3(d–f) shows different compression ratios applied to the same input scene. The parameter α controls the degree of compression: in our experiments, α is usually set to values between 0.5 and 10, higher values corresponding to a stronger compression. The application of C together with the zeroing of silhouette gradients yields modified gradient magnitudes

$$s(x,y) = \begin{cases} C(\|\nabla h(x,y)\|), & 0 \le \|\nabla h(x,y)\| < \vartheta_{\text{sil}} \, , \\ 0 & \vartheta_{\text{sil}} \le \|\nabla h(x,y)\| \end{cases} \tag{2}$$

that define slopes in the output bas-relief.

Our general depth compression approach resembles the gradient domain HDR compression by Fattal et al. [2002]. Their method applies a spatially-varying attenuation factor $\Phi(x,y)$ to the gradient magnitudes of a log-intensity image, thereby reducing large contrast and enhancing subtle contrast variations. They too fix the gradient orientation, in turn preserving the orientation of local contrast. A key property of their algorithm is to determine $\Phi(x,y)$ by combining attenuation factors $\varphi_k(x,y)$ that were determined separately

for each band of an n-frequency decomposition of the input image, with $\Phi(x,y) = \prod_k^n \varphi_k(x,y)$. Each φ_k is a non-linear compression factor in dependence on the gradient magnitude at each band, suppressing larger gradients and amplifying smaller ones. Combining the compression factors as a product across all frequencies effectively eliminates ringing, however, applied to height fields it leads to the undesirable property that geometric detail is scaled differently depending on the low frequency content of the underlying shape. Hence, we perform gradient compression on the highest resolution only. In addition, we use a different compression function that does not increase low slopes, as this would contradict the general goal of flattening the relief.

3.3 Integration

As a final step, the relief heights $h'(x,y)$ have to be recovered by integrating over the modified gradients $g'(x,y)$, which are obtained by recombining the fixed gradient orientations with the modified slopes:

$$g'(x,y) = s'(x,y)\,\mathbf{v}(x,y) \, . \tag{3}$$

However, after modifying the gradients, the vector field is not necessarily integrable anymore. Similarly to Fattal et al. [2002], we treat integration as an optimization process, finding the heights as

$$h'(x,y) = \arg\min_h \iint \|\nabla h(x,y) - g'(x,y)\|^2 \, dx\,dy \, . \tag{4}$$

This is accomplished by reformulating the optimization as a Poisson equation

$$\nabla^2 h = \operatorname{div} g' \, , \tag{5}$$

which we solve using a general multi-grid solver. Prior to integration, we downsample g' to eliminate aliasing introduced by having manipulated gradients at silhouettes.

While a unique solution to (5) exists up to a constant "sea level", it is possible to incorporate boundary conditions. We use boundary conditions wherever the boundary of the relief corresponds to its back plane. By also applying these constraints at silhouettes between scene elements and the back plane, a perfectly flat background can be enforced (cf. Figure 3(g) and Figure 7, left). For many models, however, we prefer to allow non-zero background heights, as this leads to a receding back plane along the silhouettes, emphasizing them under illumination (see Figure 6). Another way to emphasize depth discontinuities is to reinsert constant length gradients to produce "steps" where gradients have previously been zeroed (see Figure 5, bottom left, and Figure 7, right).

Due to the gradient compression, the maximum height range resulting from the integration is under (implicit) artistic control by the settings of α and ϑ_{sil}, but the exact range remains unknown until after integration. If the artist chooses an explicit height constraint, we thus linearly scale h' in order to meet this constraint, as is standard practice in tone mapping. Note that we also experimented with various constrained optimization schemes that would permit matching this range exactly *during* integration, but this leads to visible artifacts.

3.4 Artistic Controls

The presented framework already provides a powerful automated pipeline for bas-relief generation. However, depending on artistic demands, it allows for additional intuitive editing operations.

Feature Scale. A fundamental insight in sculpture is that the choice of material strongly influences the shape and composition of a relief. For instance, wood reliefs lend themselves to delicate features, while stone sculptures gain effect from strong, smoothly curved shapes. Accordingly, control over detail is a central requirement when designing a relief. We address this by including a frequency-dependent scale operator in our pipeline.

We decompose the gradient field $g'(x,y)$ into multiple frequency bands d_k, individually rescaling each band to achieve a desired detail distribution in the relief. During the frequency decomposition, special attention has to be paid to the silhouettes. We use a silhouette respecting, isotropic diffusion filter D that excludes gradient field pixels at silhouettes from the diffusion process. The linearity of D allows to perform the decomposition separately for the two components. The frequency decomposition can be described as

$$l_k = D_k(l_{k-1}), \qquad (6)$$
$$d_k = l_{k-1} - l_k, \qquad k \in \{1,\ldots,n\}, \qquad (7)$$

with $l_0 = g'$, and for a set of diffusion filters D_k of an effective radius that increases with k. Scaling each frequency band by a factor a_k and recombining them as

$$\tilde{g}' = \sum_{k=1}^{n} a_k d_k + l_n \qquad (8)$$

yields the modified gradient field $\tilde{g}'$. Figure 6 shows two different sets of scaling parameters for the dragon. On the left, low frequencies are emphasized, whereas on the right high frequencies have been amplified.

Note that this operation differs from the frequency decomposition of Fattal et al. [2002]. While in our framework, each frequency band is scaled independently, they use the decomposition to derive a common scale factor that is uniform across frequency levels.

Selective Scale. The paradigm of local, frequency-based, changes in the gradient domain is a powerful tool for further refinement. For instance, it might be desirable to emphasize or deemphasize certain scene elements by changing their apparent (depth) scale, independently of the prescale by the perspective depth buffer mapping. This can be accomplished by applying a low-frequency enhancing feature scale, confining it to gradients within a stencil that defines the scene element. This effect is demonstrated in Figure 4 wherein the image on the right contains a teapot that was emphasized further than in the unmodified version on the left. For this example, a two-band frequency decomposition was used in the feature-scale.

Figure 4: *Selective scale modification.* Left: *A bas-relief of a still life in bronze under natural environment illumination.* Right: *Scale emphasis of a scene element. Corresponding height fields are shown underneath.*

4 Results

One of our goals was to modify the geometry so that the shading of flattened relief should be close to that of the original scene. In general, however, the bas-relief may be executed in a material that differs from the original scene materials. In fact, it is common practice to make reliefs from specular materials, such as bronze or silver. We observe that more specular reliefs allow for higher depth
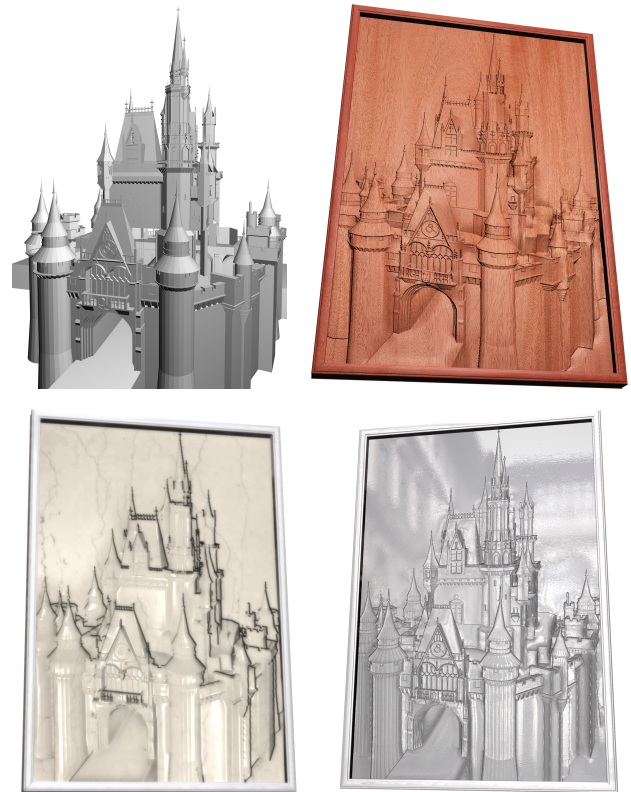

Figure 5: Left: *Original 3D model.* Right: *Three different bas-reliefs executed in wood, marble, and silver, respecting different height constraints.*

compression, as large areas of the surface have a low curvature, spreading the specular highlights to a size that resembles more diffuse lobes in an original scene. Accordingly, the final height of the relief should be chosen in correspondence with the material. Thus, the model shown in Figure 5 is mapped into three different bas-reliefs targeted at three different materials. Note that the silver relief uses a smaller range of depths, while the marble relief uses less detail. The respective ratios between panel extent and height range are 100 : 3.2 for wood, 100 : 0.88 for silver, and 100 : 4.2 for the marble relief. Another example of different height variations in conjunction with a specular material is shown in Figure 4.

We designed the table shown in relief in Figure 7, left, to challenge our algorithms. In particular, it exhibits very high-frequency silhouettes in the pencil cup and in the chains suspending the legs of the table. Furthermore, the legs have high compositional complexity with many crossings (the challenge identified in the inset figure in Section 1). One post in particular, as well as the pen case, are receding at very oblique angles, thereby challenging the range compression scheme. Nevertheless, we believe the resulting bas-relief is faithful to the medium. Finally, the vehicles and building in Figure 7, right, exhibit a high degree of depth complexity, particularly in the motorcycles, and yet the bas-relief "reads" well. Note that the truck model reveals polygonal artifacts; these are the facets of the original (low polygon count) 3D model, precisely reproduced in this medium.

To demonstrate the suitability of our process for physical media, we produced a lime stone relief, as shown in Figure 1.

5 Conclusion and Future Work

Given the generality of bas-relief, and given that digital 3D content generation increasingly pervades the creation of physical objects, we believe that automatic bas-relief generation presents an important modeling tool. We present a simple creation pipeline that is
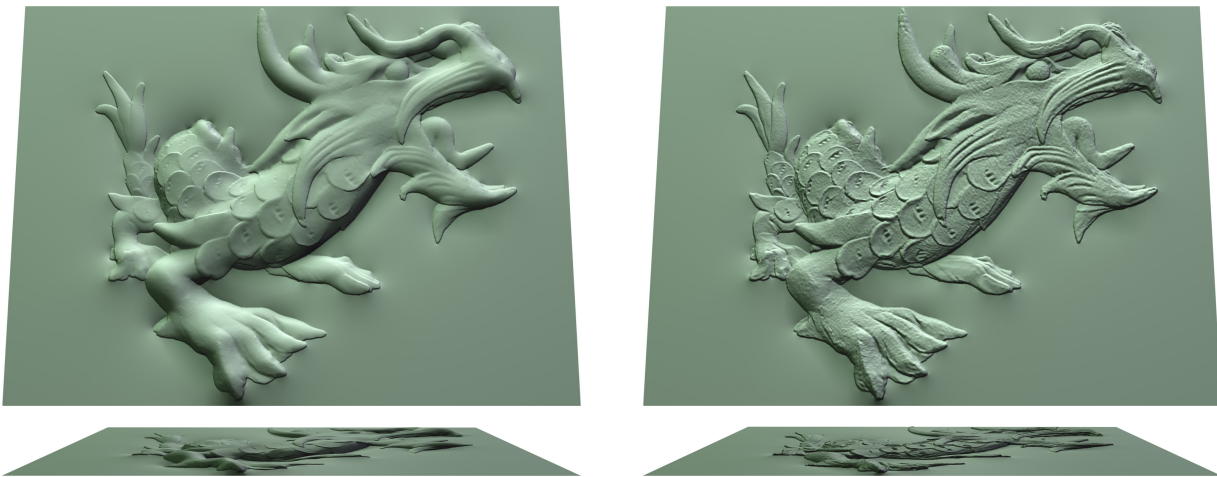
Figure 6: *Two different settings for the artistic feature scale control. Each row shows a different view of the same bas-relief.* Left: *Attenuation of fine details and increasing low frequencies.* Right: *High frequencies are amplified, while coarse shape variations are reduced. Both examples use a 3-band decomposition.*
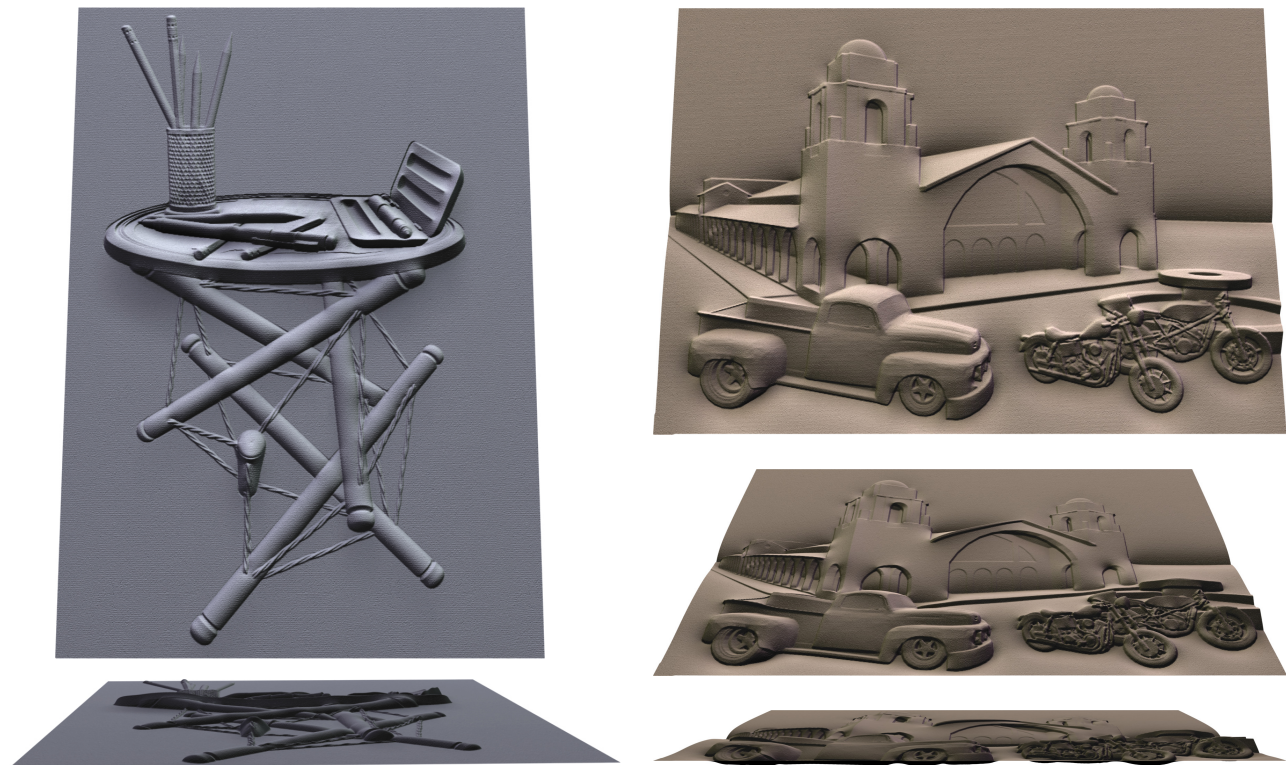

Figure 7: *Scenes with high (depth) complexity and a wide range of feature frequencies. On the right-hand side, emphasizing steps have been inserted at depth discontinuities.*

capable of producing bas-reliefs of arbitrary input scenes, at the same time meeting the various constraints imposed by the medium.

This work suggests a number of areas for future research:

Other shapes.   Although our method concentrates on the generation of bas-relief, it already addresses many of the concerns posed by creation of alto-relievo, wherein the compression requirement is somewhat relaxed. Nevertheless, other issues are unique to alto-relievo such as the potential for foreground figures to stand out in front of the background. To address the representation of this kind of object, one would probably need to abandon the height field representation in favor of a mesh. Furthermore, artists have carved re-

lief into all sorts of shapes, not just flat surfaces. We hope to explore application of relief to arbitrary surfaces, such as, for example, the decorated spheres of Yen and Séquin [2001]. The challenge would be to parameterize the relief over arbitrary shapes.

BRDF.   Material properties play a deep role in the design of bas-relief. Certain materials can sustain and reveal higher frequency detail than others. In addition, intimately connected to the "bas-relief ambiguity" is the interplay between object shape, light position, view position, and the reflectance properties of the surface [Belhumeur et al. 1999]. Our system provides controls that allow for artistic composition of the shape based on reflectance; however, it

would be of value to explore integration of material properties more deeply into the algorithmic aspects of our system in order to achieve more effective relief.

Compositing.   It is possible to imagine cases where parts of the desired scene composition are not available in an explicit geometric description but rather, e.g., as a bump or normal map, possibly obtained from real-world materials using photometric reconstruction techniques, or a distance field captured through range scanning. The intermediate gradient representation naturally lends itself to be combined with such data sources; Figure 8 shows an example relief obtained from a normal field of a real pine cone after conversion into a gradient field. Compositing different sources of data might
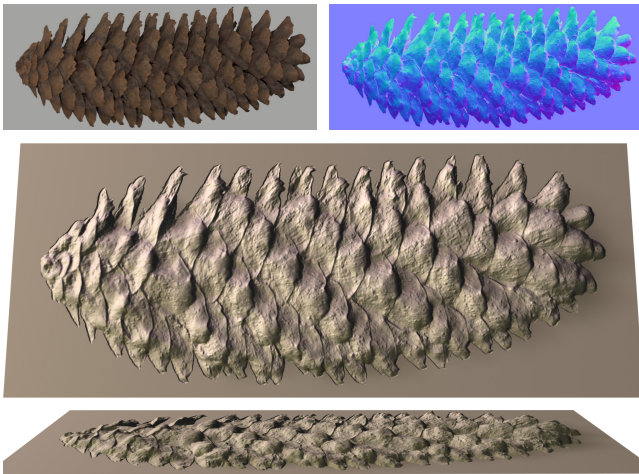

Figure 8: Top: *Photograph of a pine cone, and its normal field from photometric stereo.* Bottom: *Relief after converting normals to gradients. Gradients at silhouettes, detected as normal discontinuities, are zeroed.*

be accomplished by replacing parts of $s(x, y)$ by the gradients of an overlaid scene element, taking care to avoid large gradient discontinuities in the composition. Gradient compression and integration would then be performed on the combined gradient field.

Acknowledgments

References

ASHIKHMIN, M. 2002. A tone mapping algorithm for high contrast images. In *EGRW '02: Proceedings of the 13th Eurographics workshop on Rendering*, Eurographics Association, Aire-la-Ville, Switzerland, Switzerland, 145–156.

BELHUMEUR, P. N., KRIEGMAN, D. J., AND YUILLE, A. L. 1999. The bas-relief ambiguity. *International Journal of Computer Vision 35*, 1 (Nov.), 33–44.

CIGNONI, P., MONTANI, C., AND SCOPIGNO, R. 1997. Automatic generation of bas- and high-reliefs. *Journal of Graphics Tools 2*, 3, 15–28.

DECARLO, D., FINKELSTEIN, A., RUSINKIEWICZ, S., AND SANTELLA, A. 2003. Suggestive contours for conveying shape.

ACM Transactions on Graphics (SIGGRAPH '03) 22, 3 (July), 848–855.

DICARLO, J., AND WANDELL, B. 2000. Rendering high dynamic range images. In *Proceedings of the SPIE Electronic Imaging '2000 conference*, vol. 3965, 392–401.

DURAND, F., AND DORSEY, J. 2002. Fast bilateral filtering for the display of high-dynamic-range images. *ACM Transactions on Graphics (SIGGRAPH '02) 21*, 3 (July), 257–266.

FATTAL, R., LISCHINSKI, D., AND WERMAN, M. 2002. Gradient domain high dynamic range compression. *ACM Transactions on Graphics (SIGGRAPH '02) 21*, 3 (July), 249–256.

FLAXMAN, J. 1829. *Lectures on Sculpture*. Charles Knight, Pall Mall East, London.

GOOCH, A., GOOCH, B., SHIRLEY, P. S., AND COHEN, E. 1998. A non-photorealistic lighting model for automatic technical illustration. In *Computer Graphics (Proc. of ACM SIGGRAPH '98)*, 447–452.

HERTZMANN, A., AND ZORIN, D. 2000. Illustrating smooth surfaces. In *Computer Graphics (Proc. of ACM SIGGRAPH '00)*, 517–526.

HOCKNEY, D. 2001. *Secret Knowledge: Rediscovering the Lost Techniques of the Old Masters*. Viking Press.

HOFFMAN, M. 1939. *Sculpture Inside And Out*. W. W. Norton & Company, New York.

LARSON, G. W., RUSHMEIER, H., AND PIATKO, C. 1997. A visibility matching tone reproduction operator for high dynamic range scenes. *IEEE Transactions on Visualization and Computer Graphics 3*, 4 (October - December), 291–306.

MARKOSIAN, L., KOWALSKI, M. A., TRYCHIN, S. J., BOURDEV, L. D., GOLDSTEIN, D., AND HUGHES, J. F. 1997. Real-time nonphotorealistic rendering. In *Computer Graphics (Proc. of ACM SIGGRAPH '97)*, 415–420.

MEIER, B. J. 1996. Painterly rendering for animation. In *Computer Graphics (Proceedings of SIGGRAPH 96)*, 477–484.

RASKAR, R., ZIEGLER, R., AND WILLWACHER, T. 2002. Cartoon dioramas in motion. In *NPAR 2002: Second International Symposium on Non Photorealistic Rendering*, 7–12.

READ, H. 1961. *The Art of Sculpture*, 2nd ed. Bollingen Foundation, New York.

REINHARD, E., STARK, M., SHIRLEY, P., AND FERWERDA, J. 2002. Photographic tone reproduction for digital images. In *SIGGRAPH '02: Proceedings of the 29th annual conference on Computer graphics and interactive techniques*, ACM Press, New York, NY, USA, 267–276.

SOURIN, A. 2001. Functionally based virtual embossing. *The Visual Computer 17*, 4, 258–271.

SPÄTH, C., 2006. The digital stone project. http://digitalstoneproject.org/.

TUMBLIN, J., AND TURK, G. 1999. LCIS: A boundary hierarchy for detail-preserving contrast reduction. In *Computer Graphics (Proc. of SIGGRAPH '99)*, 83–90.

WINKENBACH, G., AND SALESIN, D. H. 1996. Rendering parametric surfaces in pen and ink. In *Computer Graphics (Proc. of ACM SIGGRAPH '96)*, 469–476.

YEN, J., AND SÉQUIN, C. 2001. Escher sphere construction kit. *2001 ACM Symposium on Interactive 3D Graphics* (March), 95–98.